

Characterization of ESnet LAN traffic at LBNL and the Comparison Between TCPDUMP Collection and NetFlow Sampling.

Esnet Measurements Team
measurements@es.net

As part of the traffic identification project, we tried to accomplish two things:

1. Characterize ESnet's LAN traffic at LBNL. The LAN serves ESnet NIC and NOC functions as well as staff offices.
2. Compare the accuracy measurements of sampled traffic (collected with NetFlow) compared with a reference total traffic capture (collected with tcpdump).

There are several potential advantages to using sampling over total packet capture for traffic identification in ESnet.

1. The volume of data collected is much smaller.
2. The collection and analysis of the data is much less hardware intensive.
3. There may be fewer collection points required.

The advantage of total packet capture is its accuracy.

This discussion assumes that the tcpdump data is authoritative.

Data Collection

Figure 1 shows the collection setup.

The data was collected for a 24-hour period from midnight on Monday, Feb 10 through midnight on Tuesday, Feb 11. The LAN traffic approximates a bell curve with a peak slightly greater than 100Mbps after midnight when system backups are done.

Because of the internal architecture of a Juniper router, Juniper recommends that the results of the data sampling should not exceed 1000 packets/sec. A sampling rate of 1:100 was chosen and worked in our environment

Simultaneous NetFlow and tcpdump packet data collections were set up to monitor ESnet's LAN traffic.

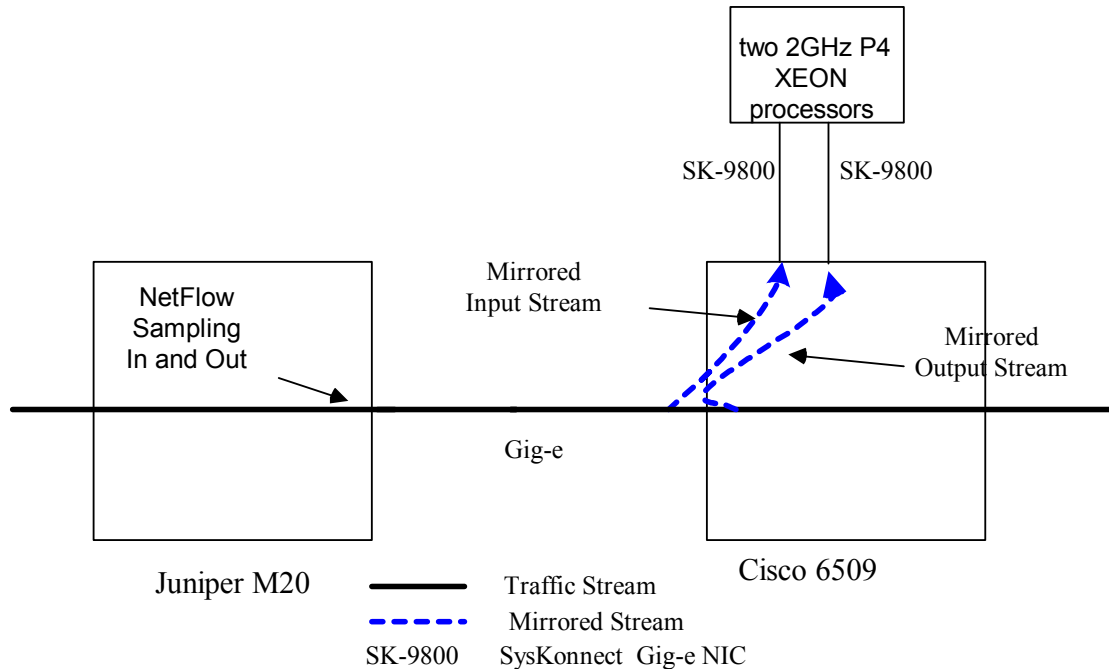


Figure 1: Diagram of the data collection set up. Input and Output sampling is enabled on the Gig-E interface between the M20 and the 6509. Mirroring of the input and output packets on the Gig-e interface in the 6509 is enabled to different Gig-e ports to a collector running tcpdump.

Table 1 describes basic characteristics of the data collected by the two methods. The smaller collection size and the smaller number of pairs sampled for the NetFlow collection is expected and is one of the advantages of using sampling over total packet capturing.

	tcpdump	NetFlow
Sampling rate	1:1	1:100
Collection size	83Gbytes	108Mbytes

Table 1: Collection statistics

Aggregation Method

Definitions:

Wk_flow All the aggregated data for a well-known port

Pair_flow All the aggregated data involving a pair of high ports (above 1000).

Ag_flow (Aggregated flows) refers to either or both of wk_flow and pair_flow.

Both sets of data were processed as follows. All the packets with either a destination port or a source port in the well know port range (0 – 1000) was collected in a bucket associated with the well know port (wk_flow). The data for these ports appear in the following tables as a destination port with no source port. The rest of the unicast data was clumped into port pairs (pair_flows), not individual flows. The raw data for the pair_flows shown in tables 3 and 5 show that they are composed of traffic between the same two hosts, but the streams may be separate in time and may not represent classic flows. For characterization of traffic and comparing the two collection methods this

aggregation method is sufficient. Since multicast traffic can only be counted on input in a Juniper router, all multicast traffic was separated from the unicast data collected by both methods. Multicast traffic collected by tcpdump was aggregated separately as described below.

Results

Table 2 shows the results of aggregating the data as described above.

	tcpdump	NetFlow	Ratio
Number ag_flows	20115	2131	9.4
Number of wk_flows	867	177	4.9
Number of pair_flows	19248	1954	9.9
Number of packets sampled	1.43E+08	1.53E+06	93.5
Number of bytes reported	1.E+11	1.15E+09	88.8

Table 2: Aggregation results

The column labeled “Ratio” is tcpdump-data/NetFlow-data. If sampling were perfect the ratio of packets sampled and bytes reported would be 100.

Table 3 shows the ag_flows with the highest byte and packet counts for both collection methods. In general the ag_flows fall in the same order for both packet and byte counts. The order is identical when comparing the highest byte counts. This is somewhat

Tcpdump top 10 Data

Sorted by Pkts						
Dst port	Src port	Pkts	Bytes	% Pkts	Cum % Pkts	Port ID
		186910509	11476361088			Multicast
20		52735180	52760553781	36.96	36.96	ftp_data
119		33867368	29557967497	23.74	60.70	nntp
161		26345566	2983871913	18.47	79.16	snmp
53		3533733	328533068	2.48	81.64	dns
49544	20013	3431232	5023323648	2.40	84.05	netflow
22		2405048	1261243718	1.69	85.73	ssh
639		2175624	277739696	1.52	87.26	msdp
57226	5155	1276577	1868908728	0.89	88.15	netflow
80		1273296	728882236	0.89	89.04	http
42990	48879	1063488	44707247	0.75	89.79	spectrum

NetFlow top 10 Data

Sorted by Pkts						
Port ID	Dst port	Src port	Pkts	Bytes	% Pkts	Cum % Pkts
Placeholder						
ftp_data	20		527644	555318442	34.58	34.578
nntp	119		380262	349814717	24.92	59.498
snmp	161		294365	41475110	19.29	78.789
netflow	49544	20013	40426	60315592	2.649	81.438
dns	53		34775	4028927	2.279	83.717
ssh	22		28008	15083139	1.835	85.552
http	80		14152	8891252	0.927	86.48
netflow	57226	5155	13991	20874572	0.917	87.396
msdp	639		12044	2042268	0.789	88.186
spectrum	42990	48879	11809	965533	0.774	88.96

Sorted by Bytes

Dst port	Src port	Pkts	Bytes	% Bytes	Cum % Bytes	Port ID
20		52735180	52760553781	51.74	51.74	ftp_data
119		33867368	29557967497	28.99	80.73	nntp
		186910509	11476361088			Multicast
49544	20013	3431232	5023323648	4.93	85.65	netflow
161		26345566	2983871913	2.93	88.58	snmp
57226	5155	1276577	1868908728	1.83	90.41	netflow
22		2405048	1261243718	1.24	91.65	ssh
56425	1969	760704	1113670656	1.09	92.74	netflow
80		1273296	728882236	0.71	93.45	http
7460	49156	327336	334934127	0.33	93.78	h323
53		3533733	328533068	0.32	94.11	dns

Sorted by Bytes

Dst port	Src port	Pkts	Bytes	% Bytes	Cum % Bytes	Port ID
ftp_data	20		527644	555318442	48.34	48.339
nntp	119		380262	349814717	30.45	78.79
Placeholder						
netflow	49544	20013	40426	60315592	5.25	84.04
snmp	161		294365	41475110	3.61	87.65
netflow	57226	5155	13991	20874572	1.817	89.468
ssh	22		28008	15083139	1.313	90.78
netflow	56425	1969	8965	13375780	1.164	91.945
http	80		14152	8891252	0.774	92.719
h323	7460	49156	4285	4493845	0.391	93.11
dns	53		34775	4028927	0.351	93.461

Table 3: A comparison of data collected by Tcpdump and NetFlow V5. Spectrum is ESnet's network management system. The high port flow labeled spectrum is between two Spectrum systems.

surprising since traffic was sampled based on packets, not bytes. The tcpdump collected multicast traffic data was aggregated separately (table 4) and the total packet. The packet and byte counts are included in the tcpdump data in the appropriate location but not used in the totals used for the % calculations.

Multicast Data		
	Pkts	Bytes Source
174888328	10728360039	Access Grid Beacon
11619601	637221147	ESnet Beacon
186507929	11365581186	Beacon Subtotal
350184	110779902	Other Multicast
186858113	11476361088	Data SubTotal
43906		PIM
8490		IGMP
52396		Protocol Subtotal
186910509	11476361088	Total Data + Protocols

Table 4: Summary of multicast data collected by tcpdump

A small amount of IPv4 encapsulated IPv6 packets are ignored.

The traffic mix is consistent with a network management location (snmp, NetFlow, Spectrum, Esnet beacon) that supplies services (nntp, http) and supports collaborative tools (h323, Access Grid Beacon, other multicast).

Table 5 further examines the traffic types and compares sampled data (NetFlow) with complete data (tcpdump). The NetFlow data through the 95%tile of bytes presented. The tcpdump data is selected to match the port pairs in the NetFlow data. The column labeled

Port ID	NetFlow V5 collection							tcpdump collection				
	Dst port	Src port	Pkts	Bytes	R	Bytes	% Cum %	Pkts	Bytes	R	Bytes	% Cum %
ftp_data	20		5.28E+05	5.55E+08	1	48.34	48.34	5.27E+07	5.28E+10	1	51.74	51.74
nntp	119		3.80E+05	3.50E+08	2	30.45	78.79	3.39E+07	2.96E+10	2	28.99	80.73
netflow	49544	20013	4.04E+04	6.03E+07	3	5.25	84.04	3.43E+06	5.02E+09	3	4.93	85.65
snmp	161		2.94E+05	4.15E+07	4	3.61	87.65	2.63E+07	2.98E+09	4	2.93	88.58
netflow	57226	5155	1.40E+04	2.09E+07	5	1.82	89.47	1.28E+06	1.87E+09	5	1.83	90.41
ssh	22		2.80E+04	1.51E+07	6	1.31	90.78	2.41E+06	1.26E+09	6	1.24	91.65
netflow	56425	1969	8.97E+03	1.34E+07	7	1.16	91.94	7.61E+05	1.11E+09	7	1.09	92.74
http	80		1.42E+04	8.89E+06	8	0.77	92.72	1.27E+06	7.29E+08	8	0.71	93.45
h323	7460	49156	4.29E+03	4.49E+06	9	0.39	93.11	3.27E+05	3.35E+08	9	0.33	93.78
dns	53		3.48E+04	4.03E+06	10	0.35	93.46	3.53E+06	3.29E+08	10	0.32	94.11
h323	7462	49204	4.02E+03	3.92E+06	11	0.34	93.80	2.31E+05	2.18E+08	13	0.21	94.86
h323	49156	7466	4.06E+03	3.86E+06	12	0.34	94.14	2.30E+05	2.16E+08	15	0.21	95.28
h323	7464	3232	3.30E+03	3.49E+06	13	0.30	94.44	2.12E+05	2.16E+08	14	0.21	95.07
ms-ds	445		8.78E+03	3.08E+06	14	0.27	94.71	8.82E+05	2.71E+08	12	0.27	94.64
h323	49162	7462	2.31E+03	2.25E+06	15	0.20	94.90	2.26E+05	2.14E+08	16	0.21	95.49
h323	7462	49162	2.23E+03	2.18E+06	16	0.19	95.09	2.27E+05	2.14E+08	17	0.21	95.70
msdp	639		1.20E+04	2.04E+06	17	0.18	95.27	2.18E+06	2.78E+08	11	0.27	94.38
imap/ssl	993		2.97E+03	1.69E+06	18	0.15	95.42	1.77E+05	8.82E+07	32	0.09	97.39
h323	49156	7460	2.17E+03	1.68E+06	19	0.15	95.57	1.63E+05	1.20E+08	22	0.12	96.37
netflow	53377	20229	1.06E+03	1.58E+06	20	0.14	95.70	8.85E+04	1.30E+08	21	0.13	96.25
smtp	25		7.25E+03	1.54E+06	21	0.13	95.84	6.11E+05	1.08E+08	27	0.11	96.93

Table 5: This table matches ports through the 95%tile (based on bytes) NetFlow data with the corresponding data from the tcpdump collection.

The columns labeled "R" is the rank in the respective collection based on descending byte counts,

Rnk is the rank of the port pair's data relative to the other data in the collection. For the first 10 pairs there is good agreement between the ranks of the ag_flows. After rank 10, the divergence in rank between matched ag_flows increases.

Table 6 shows a more detailed comparison between the NetFlow and tcpdump collected data. NetFlow reports packet size while tcpdump reports the amount of data in a packet. The column labeled "Adj Bytes" adjusts the byte count for the transport protocol header length. For each TCP packet (marked by a "T" in the P column) 40 bytes/packet are added to the byte total. For each UDP packet (marked by a "U" in the P columns) 28 bytes are added to the byte total. The columns labeled TCP/NetFlow are the ratio (tcpdump collected data)/(sampled data). The comparison between packet and byte totals is shown. If we had perfect sampling, we would expect ratio to be 100. The packet count ratios vary by large amounts, but the first 10 seem better than the remainder. Two columns are shown for the byte ratios. The first is calculated using the reported byte counts. The second (Bytes Fixed) uses the "Adj Bytes". Adjusting the byte counts for the header size does significantly impact the ratio. Once again, the first 10 ratios are generally closer to 100 than the remainder.

The Packet Size column shows the average size of the packets collected by the two methods for each of the reported ag_flows. The corrected byte totals are used for the tcpdump data. The agreement is very good and indicates the data collected by the two methods is the same, and that the differences are most likely due to errors inherent in sampling small ag_flows rather than NetFlow over or under reporting certain packet types or sizes.

Port ID	NetFlow V5 collection					tcpdump collection					Pkts		Bytes			Packet size		
	Dst	Src	Pkts	Bytes	R	Pkts	Bytes	P	Adj. Bytes	R	TCP/	%	TCP/	%	Fixed			
	port	port									Netflw	diff	Netflw	diff				
ftp_data	20		5.28E+05	5.55E+08	1	5.27E+07	5.28E+10	T	5.49E+10	1	99.94	0.06	95.01	4.99	98.81	1052	1040	0.99
nntp	119		3.80E+05	3.50E+08	2	3.39E+07	2.96E+10	T	3.09E+10	2	89.06	10.94	84.50	15.50	88.37	920	913	0.99
netflow	49544	20013	4.04E+04	6.03E+07	3	3.43E+06	5.02E+09	U	5.12E+09	3	84.88	15.12	83.28	16.72	84.88	1492	1492	1.00
snmp	161		2.94E+05	4.15E+07	4	2.63E+07	2.98E+09	U	3.72E+09	4	89.50	10.50	71.94	28.06	89.73	141	141	1.00
netflow	57226	5155	1.40E+04	2.09E+07	5	1.28E+06	1.87E+09	U	1.90E+09	5	91.24	8.76	89.53	10.47	91.24	1492	1492	1.00
ssh	22		2.80E+04	1.51E+07	6	2.41E+06	1.26E+09	T	1.36E+09	6	85.87	14.13	83.62	16.38	90.00	539	564	1.05
netflow	56425	1969	8.97E+03	1.34E+07	7	7.61E+05	1.11E+09	U	1.13E+09	7	84.85	15.15	83.26	16.74	84.85	1492	1492	1.00
http	80		1.42E+04	8.89E+06	8	1.27E+06	7.29E+08	U	7.65E+08	8	89.97	10.03	81.98	18.02	85.99	628	600	0.96
h323	7460	49156	4.29E+03	4.49E+06	9	3.27E+05	3.35E+08	U	3.44E+08	9	76.39	23.61	74.53	25.47	76.57	1049	1051	1.00
dns	53		3.48E+04	4.03E+06	10	3.53E+06	3.29E+08	U	4.27E+08	10	101.62	-1.62	81.54	18.46	106.10	116	121	1.04
h323	7462	49204	4.02E+03	3.92E+06	11	2.31E+05	2.18E+08	U	2.24E+08	13	57.56	42.44	55.56	44.44	57.21	975	970	0.99
h323	49156	7466	4.06E+03	3.86E+06	12	2.30E+05	2.16E+08	U	2.23E+08	15	56.56	43.44	56.08	43.92	57.75	949	969	1.02
h323	7464	3232	3.30E+03	3.49E+06	13	2.12E+05	2.16E+08	U	2.22E+08	14	64.09	35.91	61.96	38.04	63.66	1057	1050	0.99
ms-ds	445		8.78E+03	3.08E+06	14	8.82E+05	2.71E+08	U	2.96E+08	12	100.46	-0.46	88.08	11.92	96.11	350	335	0.96
h323	49162	7462	2.31E+03	2.25E+06	15	2.26E+05	2.14E+08	U	2.20E+08	16	97.90	2.10	95.03	4.97	97.84	973	972	1.00
h323	7462	49162	2.23E+03	2.18E+06	16	2.27E+05	2.14E+08	U	2.20E+08	17	102.11	-2.11	97.99	2.01	100.91	979	968	0.99
msdp	639		1.20E+04	2.04E+06	17	2.18E+06	2.78E+08	T	3.65E+08	11	180.64	-80.64	136.00	-36.00	178.61	170	168	0.99
imap/ssl	993		2.97E+03	1.69E+06	18	1.77E+05	8.82E+07	T	9.53E+07	32	59.70	40.30	52.13	47.87	56.32	570	538	0.94
h323	49156	7460	2.17E+03	1.68E+06	19	1.63E+05	1.20E+08	U	1.25E+08	22	75.22	24.78	71.44	28.56	74.16	777	766	0.99
netflow	53377	20229	1.06E+03	1.58E+06	20	8.85E+04	1.30E+08	U	1.32E+08	21	83.30	16.70	81.74	18.26	83.30	1492	1492	1.00
smtp	25		7.25E+03	1.54E+06	21	6.11E+05	1.08E+08	T	1.32E+08	27	84.21	15.79	70.15	29.85	86.04	212	203	0.96

Table 6: Further comparison of packet sampling vs total packet capture

This table matches ports through the 95%tile (based on bytes) NetFlow data with the corresponding data from the tcpdump collection.

P is the underlying transport protocol (T=tcp, U=UDP)

Adj Bytes includes the IP and Protocol header bytes

R is the rank in the respective collection based on descending byte counts,

TCP/Netflw columns are the ratio (tcpdump data)/(netflow data) using packet and byte counts,.

Packet size is the average packet size for NetFlow and tcpdump collected data.

Ratio is the (packet size from tcpdump data)/packet size from NetFlow data)

In summary, ESnet LAN traffic was collected and analyzed by two methods:

Sampling of packets at a rate of 1:100 and collecting the data with NetFlow.
Total capture using tcpdump.

The traffic mix was found to be consistent with a location that has significant network management functions, supplies services to the community and supports collaborative tools.

The accuracy of scaling sampled data to represent a full collection was examined. If sampling were 100% perfect, the ratio between NetFlow and tcpdump would be 100, reflecting the sampling rate of 1:100. The accuracy of scaling, as reflected by the ratio, tends to decrease as the ag_flow size decreases. For some of the smaller ag_flows, the ratio drops to almost 50. The excellent agreement of packet size between the two methods indicates that the difference is due errors inherent in sampling smaller ag_flows. In the test case presented here, NetFlow ag_flows accounting for up to 92% of the total bytes, are within 15% of the expected (tcpdump) value. If an inaccuracy of 10 to 15% is acceptable, sampling could be used to estimate large flows. For example, the large bucket for ftp_data scales well, and nntp is not too bad. Those two buckets account for 80% of our traffic.

At a site with more traffic, the sampling rate would need to decrease. For a Gig-e connected site we have used a rate of 1:500. Higher speed connections would require a further decrease in the sampling rate. This would increase the error in extrapolating from sampled data to full data at larger site.